CD-1.1 Reliable Multicasting: Status and Plans

Ray Cordova, SAIC
Roger Bowman, SAIC
Deb Agarwal, LBNL
Gus Gustafson, Consultant

October 22, 2002

Motivation for CD-1.1 Reliable Multicast

- Enhance reliability of IMS data transmission to signatories
- Minimize reliability and data distribution burden (= Cost) for the IDC
- Provide direct transmission of IMS station data to NDC hosting station and IDC
- Minimize bandwidth required for transmission of data from a single data producer to multiple destinations

What is Multicast Communication?

- Group communication mechanism
  - Provides one-to-many and many-to-many communication
- Efficient dissemination of messages
  - Network-based duplication (when needed)
  - Multicast retransmissions
  - Bandwidth savings
  - Parallel delivery at multiple locations

IP Multicast Communication

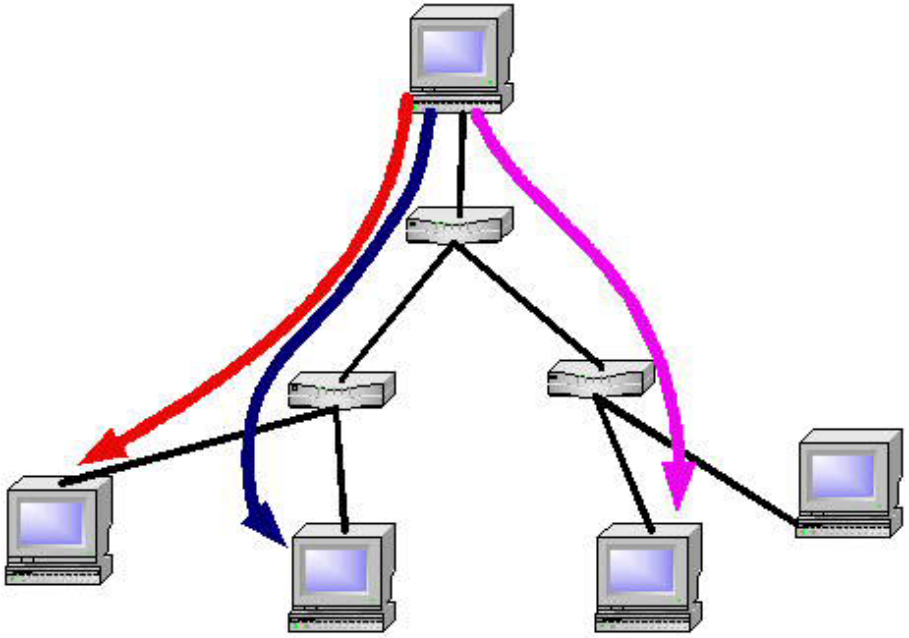

Unicast

Multicast

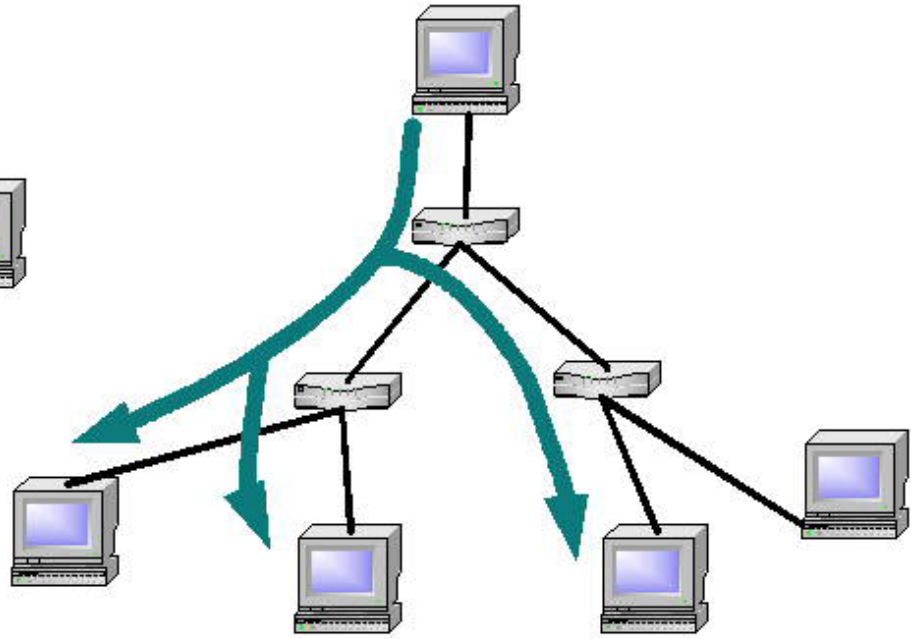

Example IP Multicast Use (Access Grid)

What is Reliable Multicast?

- Properties similar to TCP
- Application-level program
- Uses IP Multicast as the underlying communication mechanism
- Reliable and ordered delivery of messages within a group
- Tracks group membership
- IETF Reliable Multicast Transport Working Group is defining standardized building blocks

Example Reliable Multicast Use (Remote Instrument Access)

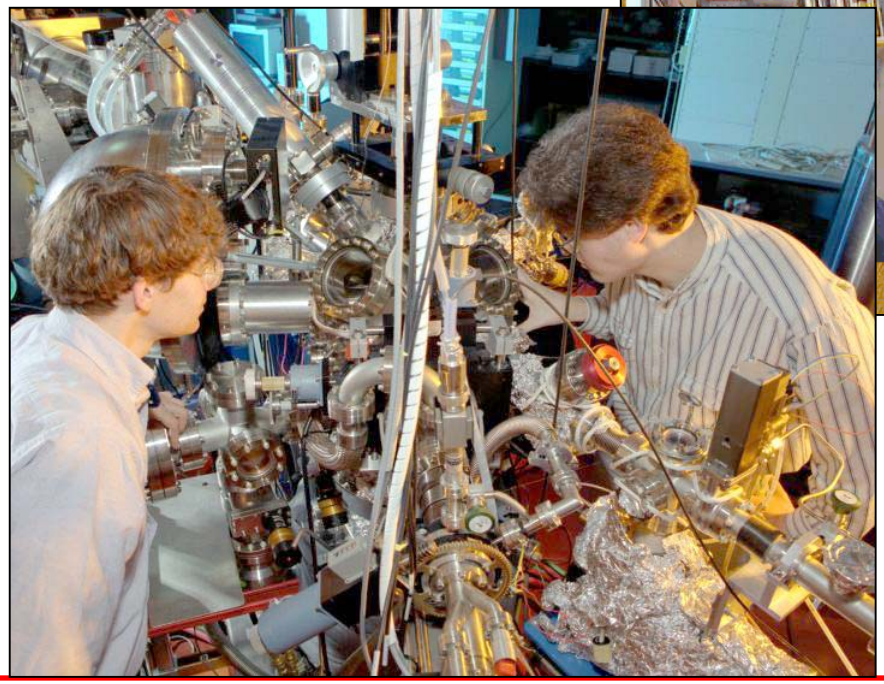

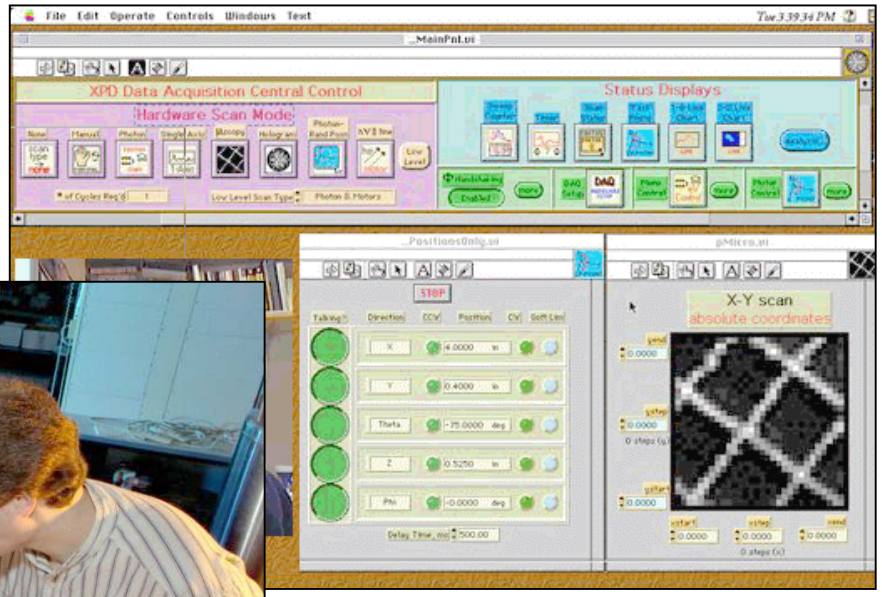

CD-1.1 Reliable Multicast

- Assumptions
  - Reliability of CD-1.1 multicasting should be equivalent to that of CD-1.1 unicasting
  - Bandwidth of tail circuit to stations will not be increased to support multicast (in catch-up mode)

System-Level Requirements

- All system-level CD-1.1 unicast requirements apply
  - 99.99% reliability requirement has been retained
    - Note that this was tightened from 99.9% in [IDC7.4.1]
    - Application level reliability provided by a combination of multicast and point-to-point transmission mechanisms
  - The connection requirement has been modified to generalize initiation of a connection after an outage
- New system-level requirement
  - “The CDS shall provide multicasting capability to deliver CD-1.1 frames using underlying IP router multicast capability.”

Design Constraints

- Minimize perturbations to CD-1.1 Formats and Protocols [IDC 3.4.3 Rev. 0.2]
- Compatibility with CDS CD-1.1 unicast
  - A data provider must be able to service both multicast and point-to-point data consumers simultaneously
- CDS CD-1.1 software components to operate in either multicast or unicast mode, according to parameter settings

Design Approach

- System level
  - Use custom reliable multicast solution (based on CD1.1)
  - Use point-to-point mechanism to provide application level reliability
  - Separate multicast (real-time) and point-to-point (catch-up) transmission into separate subsystems
- Multicast subsystem
  - Multicast transmission initiated by pull from data consumer
  - Use CD-1.1 procedure (in reverse) to establish connection
  - Multicast transmission begins at “data provider time” (current time - small look back) - no attempt to catch up
  - Data provider provides packet-level reliability host

CD-1.1 Reliable Multicast Participants
An Example - Normal Operation

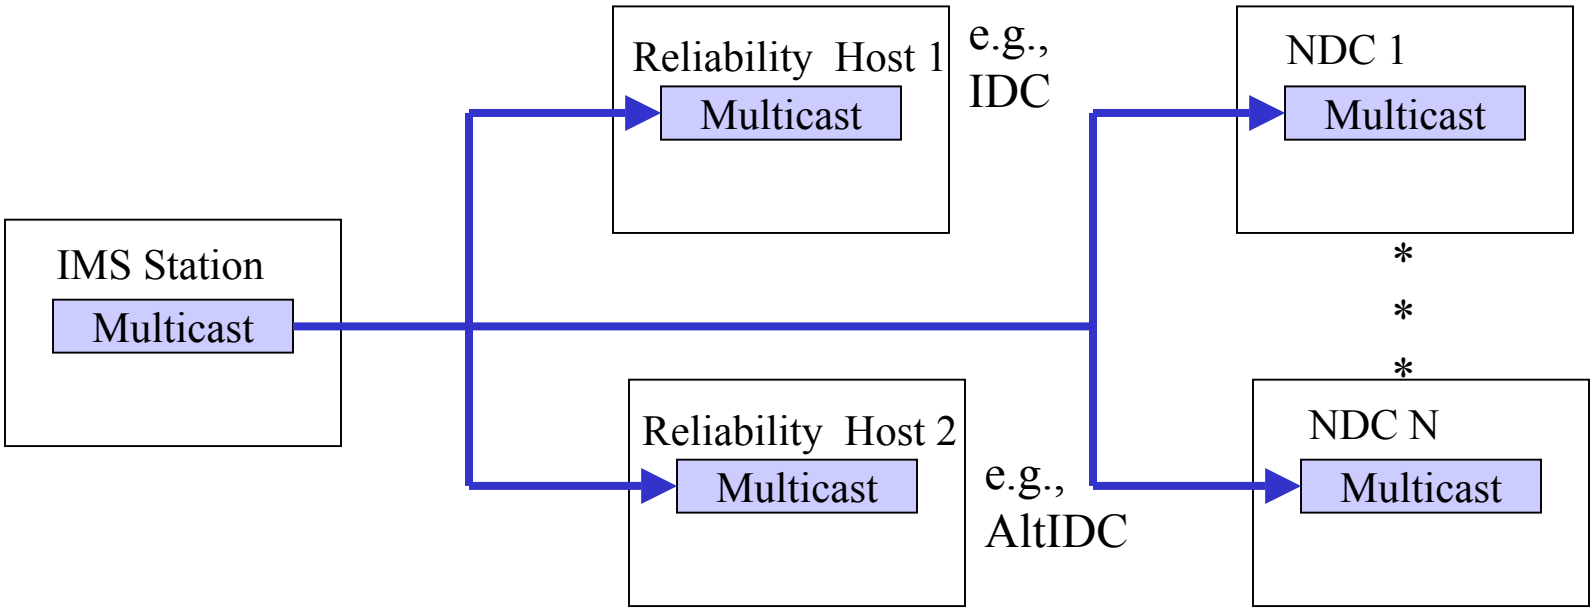

CD-1.1 Reliable Multicast Participants
An Example - Catch-Up Operation

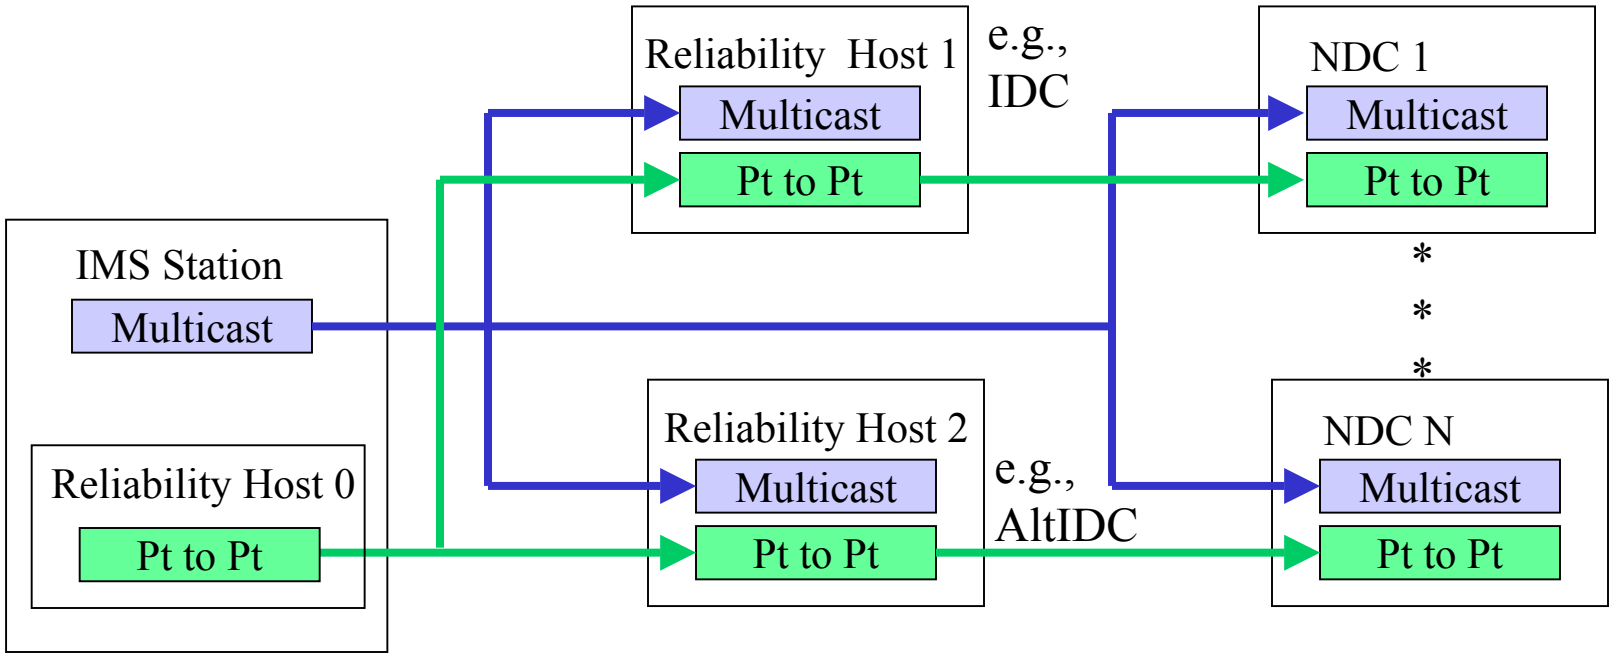

Conceptual Design

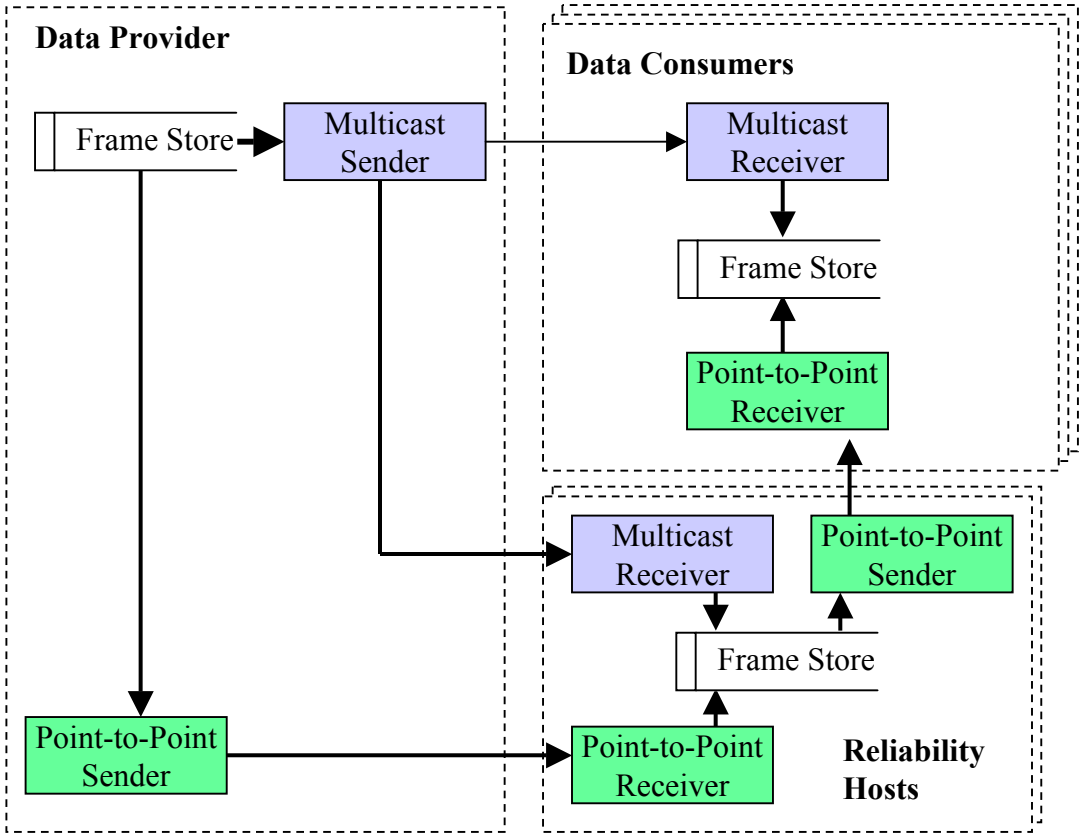

Response to Failure Modes

Failure mode	Design response
Packet loss (normal mode)	1. Consumer sends PNack(s). 2. Provider re-multicasts missing packet(s).
Network congestion (normal mode)	1. Provider is limited to a maximum UDP transmission rate, which will be set to account for predicted TCP traffic.
Station or tail circuit outage	1. All multicast sessions disconnect. 2. Consumers initiate multicast reconnections with Provider. 3. Provider begins to multicast at current time minus small lookback. 4. If outage duration exceeds lookback, Consumer identifies missing Data Frames and initiates point-to-point reconnection with Reliability Host. 5. Consumer pulls missing Data Frames from Reliability Host.

Response to Failure Modes (2)

Failure mode	Design response
Reliability Host outage	No consequence to Consumers if multicast and packet reliability service from Provider operates without error. If a Data Frame has been determined to be missing, then: 1. Consumers keep trying to connect with Reliability Host until successful. 2. Reliability Host connects to Provider and pulls data using point-to-point mechanism. 3. Consumers pull missing Data Frames from Reliability Host.
Regional network node outage	1. Multicast sessions with affected Consumers disconnect. 2. Affected Consumers initiate multicast reconnection with Provider. 3. Affected Consumers initiate point-to-point reconnection with their Reliability Hosts. 4. Consumers pull missing Data Frames from Reliability Hosts.

Response to Failure Modes (3)

Failure mode	Design response
Data Consumer or tail circuit outage	1. Multicast session with Provider disconnects. 2. Consumer initiates multicast reconnection with Provider. 3. Consumer initiates point-to-point reconnection with its Reliability Host. 4. Consumer pulls missing Data Frames from Reliability Host.

Plans

- Finalize design
  - Examine failure modes and design response
  - Conduct design review
  - Update CD-1.1 Formats and Protocols document for multicast design
- Implement new multicast software components
- Adapt existing components of Continuous Data Subsystem CD-1.1, as necessary

Plans (2)

- Plan and conduct wide-area multicast testing
  - Phase 1:
    - Data provider at SAIC Testbed (San Diego)
    - Data consumers at CMR and SAIC Testbed
    - First quarter of 2003
  - Phase 2:
    - Data provider at I57US
    - Data consumers at CMR and SAIC Testbed or US NDC
    - Second quarter of 2003
  - Possible Phase 3:
    - Partitioned Subnetwork Configuration -
    - IMS multicast station to host NDC and IDC

Design Issues

- Pull model for multicast transmission connection
  - (Pull model for point-to-point is not an issue)
- Negotiation/publication of Formats & Protocols revisions